First Class Futures: a Study of Update Strategies

Muhammad Uzair Khan, Ludovic Henrio

▶ **To cite this version:**

Muhammad Uzair Khan, Ludovic Henrio. First Class Futures: a Study of Update Strategies. [Research Report] RR-7113, INRIA. 2009. ⟨inria-00435573⟩

HAL Id: inria-00435573

https://inria.hal.science/inria-00435573v1

Submitted on 24 Nov 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

HAL Authorization

INSTITUT NATIONAL DE RECHERCHE EN INFORMATIQUE ET EN AUTOMATIQUE

First-class Futures: a Study of Update Strategies

Khan Muhammad — Ludovic Henrio

N° 7113

November 2009

ISSN 0249-6399 ISRN INRIA/RR--7113--FR+ENG

INSTITUT NATIONAL
DE RECHERCHE
EN INFORMATIQUE
ET EN AUTOMATIQUE

INRIA

centre de recherche
SOPHIA ANTIPOLIS - MÉDITERRANÉE

First-class Futures: a Study of Update Strategies

Khan Muhammad * , Ludovic Henrio *

Thème : Distributed systems
Équipes-Projets Oasis

Rapport de recherche n° 7113 — November 2009 — 14 pages

Abstract: Futures enable an efficient and easy to use programming paradigm for distributed applications. A natural way to benefit from distribution is to perform asynchronous invocations to methods or services. Upon invocation, a request is en-queued at the destination side and the caller can continue its execution. But a question remains: "what if one wants to manipulate the result of an asynchronous invocation?" First-class futures provide a transparent and easy-to-program answer: a future acts as the placeholder for the result of an asynchronous invocation and can be safely transmitted between processes while its result is not needed. Synchronization occurs automatically upon an access requiring the result. As references to futures disseminate, a strategy is necessary to propagate the result of each request to the processes that need it. This report studies the efficient transmission of results: it presents three main strategies in a semi-formal manner, and provides a cost analysis with some experiments to determine the efficiency of each strategy.

Key-words: distributed systems, future, future update strategies, active objects, asynchronous communications

* INRIA, Sophia Antipolis, Université de Nice – I3S – CNRS

Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles, BP 93, 06902 Sophia Antipolis Cedex
Téléphone : +33 4 92 38 77 77 — Télécopie : +33 4 92 38 77 65

Futurs de primière classe: Une étude des stratégies de mise à jour des futures

Résumé : Une futur est un paradigme de programmation efficace et facile à utiliser pour les applications distribuées. Un futur est une espace réservé temporaire pour un résultat d'appel asynchrone. Cet article étudie la transmission efficace des résultats: Elle présente trois stratégies demanière semi-formelle, et fournit une analyse des coûts avec quelques expériences de manière à déterminer l'efficacité de chaque stratégie.

Mots-clés : systèmes distribués, future, stratégies pour mise à jour les futures, objet active, communication asynchrone

Contents

1 Introduction **4**

2 Related works **4**

3 Modeling Different Future Update Strategies **5**
3.1 General Notation 5
3.1.1 Operations 6
3.1.2 Events 6
3.2 Eager Forward-Based Strategy 7
3.2.1 Send Future Reference 7
3.2.2 Future Computed 8
3.2.3 Send Future Value 8
3.3 Eager Message-based Strategy 8
3.3.1 Send Future Reference 8
3.3.2 Future Computed 9
3.3.3 Send Future Value 9
3.4 Lazy Message-based Strategy 9
3.4.1 Send future reference 10
3.4.2 Wait-by necessity 10
3.4.3 Future Computed 10
3.4.4 Send Future Value 10
3.5 Cost Analysis of Update Strategies 10

4 Experimental Evaluation **11**

5 Conclusion **13**

1 Introduction

Futures are language constructs that improve concurrency in a natural and transparent way. A *future* is used as a place holder for a result of a concurrent computation [7, 16]. Once the computation is complete and a result (called *future value*) is available, the placeholder is *replaced* by the result. Access to an unresolved future is a blocking operation. As results are only awaited when they are really needed, computation is parallelized in a somehow optimal way. The future creation can be *transparent* or *explicit*. With explicit futures, specific language constructs are necessary to create the futures and to fetch the result. *Transparent* futures, on the other hand, are managed by the underlying middleware and the program syntax remains unchanged; futures have the same type as the actual result. Some frameworks allow futures to be passed to other processes. Such futures are called *First class futures* [3]. In this case additional mechanisms to update futures are required not only on the creating node, but also on all nodes that receive a future. First class futures offer greater flexibility in application design and can significantly improve concurrency both in object-oriented and procedural paradigms like workflows [14, 13]. They are particularly useful in some design patterns for concurrency, such as master-worker and pipeline.

Our work analyzes several future update strategies; it can be considered as an extension of [3] and [11] through a language-independent approach that makes it applicable to various existing frameworks that support first class futures. The experiments are performed with the ProActive library [1], which is a middleware providing first-class futures. Our main contributions are: a semi-formal event-like notation to model the future update strategies, and a description of three different update strategies using this notation (Section 3); a cost-analysis of the presented protocols (Section 3.5); some experimental results carried out with the ProActive library (Section 4).

2 Related works

Futures, first introduced in Multilisp [7] and ABCL/1 [16] are used as constructs for concurrency and data flow synchronization. Languages/frameworks that make use of explicit constructs for creating futures include Multilisp [7, 6], λ-calculus [10], SafeFuture API [15] and ABCL/f [12]. In contrast, futures are created implicitly in frameworks like ASP [3, 2], AmbientTalk [5] and ProActive [1]. This implicit creation corresponds to asynchronous invocation. A key benefit of the implicit creation is that no distinction is made between local and remote operations in the program. Additionally, the futures can be accessed explicitly or implicitly. In case of explicit access, operations like *claim* and *get*, *touch* are used to access the future [9, 8, 12]. For implicit access, operations that need the real value of an object are called *strict operations*, and automatically trigger an access to the future. Accessing a future that has not been updated, results in a *wait-by-necessity*, which blocks the accessing process.

Creol [8] allows for explicit control over data-flow synchronizations. In [4], Creol has been extended to support first class futures even if the access is still explicit (using *get* and *await*). ASP [3] and ProActive [1], have transparent first-class futures. Thus, the synchronization is transparent and data-flow oriented.

In AmbientTalk, futures are also first-class and transparently manipulated; but the future access is a non-blocking operation: it is an asynchronous call that returns another future. This avoids the possibility of a dead lock as there is no synchronization. In AmbientTalk a future contains a *mailbox* to store the received invocations, which are treated upon future update. This significantly differs from the approach adopted in frameworks like [8] where access to a future is blocking. All processes interested in the future are registered as observers. When the result for the future is computed, all the registered observers are notified, thus the future update strategy in AmbientTalk is closed to the eager-message based strategy presented here. [15] provides a *safe* extension to Java futures, but with explicit creation and access.

This report presents future update strategies, with more details than in [3]; it also analyzes the cost of each strategy. This work is language independent, and can be applied to any frameworks supporting first class futures, e.g. [4, 5].

3 Modeling Different Future Update Strategies

This section gives a semi-formal definition for the three main future update strategies. Strategies are called *eager* when all the references to a future are updated as soon as the future value is calculated. They are called *lazy* if futures are only updated upon need, which minimizes communications but might increase the time spent waiting for the future value. Two eager and one lazy strategies are presented here: *eager forward-based* (following the future flow), *eager message-based* (using a registration mechanism, also called home-based in [11]), and *lazy message based*. One could also consider a lazy forward-based strategy, but as it is extremely inefficient, we do not discuss it here.

3.1 General Notation

This section presents a brief overview of the various notation and entities that we use to model the future update strategies. We denote by $\mathcal{A}$ the set of processes (also called *activities*); α, β, $\ldots \in \mathcal{A}$ range over processes. $\mathcal{F}$ denotes the set of future identifiers, each future identifier is of the form $f^{\alpha\rightarrow\beta}$, which represents the future f created by the activity α, and being calculated by β. As each object needs to keep track of the futures it has received, we make use of some local lists for this purpose. There is one *future list* for each activity α. It represents the location where the futures are stored in local memory.

$$\mathcal{FL}_\alpha \colon \mathcal{F} \mapsto \mathcal{P}(Loc)$$

Locations, called *loc* in the following and of type *Loc*, refer to the in-memory position of the future. To keep track of activities to which a future is to be sent, a *future recipient* list is stored in each process.

$$\mathcal{FR}_\delta \colon \mathcal{F} \mapsto \mathcal{P}(\mathcal{A})$$

$\gamma \in \mathcal{FR}_\delta\,(f^{\alpha\rightarrow\beta})$ if the future value for $f^{\alpha\rightarrow\beta}$ has to be sent from δ to γ. It should be noted that each $f^{\alpha\rightarrow\beta}$ can be mapped to several locations in $\mathcal{FL}$ or several activities in $\mathcal{FR}$. $\mathcal{FR}$ and $\mathcal{FL}$ are initialized to empty mapping on all processes. We use an event like notation to define the different strategies. Operations triggered by the strategies, and events triggered by the rest of the middleware are described respectively in bellow. Events are indexed by the activity on which they occur, or $\alpha \rightarrow \beta$ for a communication from α to β.

3.1.1 Operations

Register Future - Reg: $\mathcal{F} \times \mathcal{B} \times \mathcal{F} \mapsto \mathcal{P}(\mathcal{B})$
We define an operation *Reg* that is given a future, a process and a mapping $\mathcal{F} \mapsto \mathcal{P}(\mathcal{B})$ (either $\mathcal{FL}$ when $\mathcal{B} = Loc$, or $\mathcal{FR}$ when $\mathcal{B} = \mathcal{A}$). $Reg_\gamma(f^{\alpha\to\beta}, b, L)$ replaces the list L by the list L' defined as follows:

$$L'(f_2^{\alpha'\to\beta'}) = \begin{cases} L(f^{\alpha\to\beta}) \cup \{b\} & \text{if } f_2^{\alpha'\to\beta'} = f^{\alpha\to\beta} \\ L(f_2^{\alpha'\to\beta'}) & \text{else} \end{cases}$$

The *Reg* operation replaces the old mapping L with a new one containing the additional mapping. An example of its usage could be $Reg_\gamma(f^{\alpha\to\beta}, loc, \mathcal{FL}_\gamma)$ which adds to the $\mathcal{FL}_\gamma$ list, a new location *loc* associated to future $f^{\alpha\to\beta}$.

Locally Update future with value - Update: $Loc \times Value$
Once the value for a given future is received, this operation is triggered to update all corresponding local futures with this value. The operation $Update_\gamma(f^{\alpha\to\beta}, v)$ replaces, in the activity γ, each reference to the future $f^{\alpha\to\beta}$ by the value v. Remember the set of locations of these references is $\mathcal{FL}_\gamma(f^{\alpha\to\beta})$.

Clear future from list - Clear: $\mathcal{F} \times \mathcal{F} \mapsto \mathcal{P}(\mathcal{B})$
The clear operation $Clear(f^{\alpha\to\beta}, L)$ removes the entry for future $f^{\alpha\to\beta}$ from the list L; it replaces the list L by the list L' defined by:

$$L'(f_2^{\alpha'\to\beta'}) = \begin{cases} L(f_2^{\alpha'\to\beta'}) & \text{if } f_2^{\alpha'\to\beta'} \neq f^{\alpha\to\beta} \\ \emptyset & \text{else} \end{cases}$$

It will be used after a future update to clear entries for the updated future.

Send future value: SendValue: $\mathcal{F} \times Loc \times Value$
Send operation is used when a process needs to send the value of a computed future to another process in order to update the future there. $SendValue_{\delta\to\gamma}(f^{\alpha\to\beta}, loc, v)$ sends the value v for the future $f^{\alpha\to\beta}$ from δ to γ. Sending a future value can trigger send future reference events, $SendRef$, for all the future references contained in the value v. The exact details of this operation depend on the strategy and appear in Sections 3.2, 3.3, and 3.4

3.1.2 Events

Future update strategies react to events, triggered by the application or the middleware, presented below.

Create future: Create: $\mathcal{F} \times Loc$
$Create_\alpha(f^{\alpha\to\beta}, loc)$ is triggered when α creates a future that will be calculated by the process β. The semantics of this event is similar for all strategies: it registers the future in the future list $\mathcal{FL}$ of the creating process.

$$Create_\alpha(f^{\alpha\to\beta}, loc) \triangleq Reg_\alpha(f^{\alpha\to\beta}, loc, \mathcal{FL}_\alpha)$$

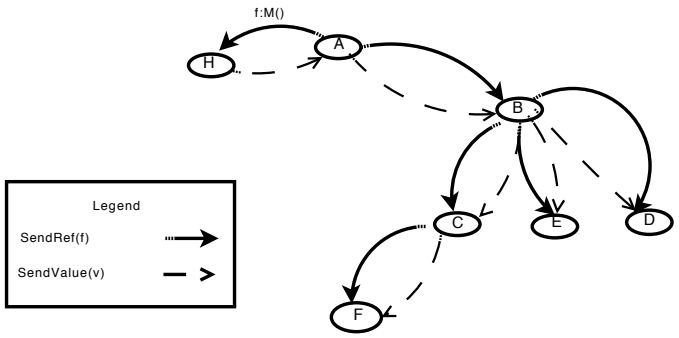

Figure 1: Future-update in eager forward-based strategy

Send future reference: SendRef: $\mathcal{F} \times Loc$
$SendRef_{\delta \to \gamma}(f^{\alpha \to \beta}, loc)$ occurs when the process δ sends the future reference $f^{\alpha \to \beta}$ to γ and the future is stored at the location *loc* on the receiver side. The exact details of this operation depend on the strategy and will be described in Sections 3.2, 3.3, and 3.4.

Future computed: FutureComputed: $\mathcal{F} \times Value$
$FutureComputed_{\beta}(f^{\alpha \to \beta}, val)$ occurs when the value *val* of future $f^{\alpha \to \beta}$ has been computed by β. Reactions to this event will be described below.

Wait-by-necessity: Wait: $\mathcal{A}$
This event is triggered when a process accesses an unresolved future. This corresponds to *get* or *touch* operation in [9, 8, 12]. For the two eager strategies it simply causes the process to be blocked until the value is received. For the lazy strategy, this event retrieves the future value, see Section 3.4.

3.2 Eager Forward-Based Strategy

In this strategy, each process remembers the nodes to which it has forwarded the future. When the value is available, it is sent to all such nodes. The list of processes to which a process β should send the future value for $f^{\alpha \to \beta}$ is $\mathcal{FR}_{\beta}(f^{\alpha \to \beta})$. It is the list of processes to which β has sent the future reference.

Figure 1 shows an example illustrating this strategy. Process A makes an asynchronous call on process H and receives the future $f^{A \to H}$. A then passes this future to B, which in turn passes the future to C, D and E. Finally C passes the future to F. Each time a future is forwarded, i.e. upon a *SendRef* message, the forwarding process δ adds the destination to its $\mathcal{FR}_{\delta}(f^{A \to H})$. When the result for $f^{A \to H}$ is available, it is communicated to A using *SendValue* message. A then forwards the update on B ($\mathcal{FR}_A(f^{A \to H}) = \{B\}$). B can make concurrent updates on C, E and D ($\mathcal{FR}_B(f^{A \to H}) = \{C, E, D\}$). Finally, the occurrence in F is updated by C ($\mathcal{FR}_F(f^{A \to H}) = \{C\}$).

3.2.1 Send Future Reference

When a process δ sends a future $f^{\alpha \to \beta}$ to a process γ, the sender registers the destination process in $\mathcal{FR}_{\delta}$, and the destination process registers the location

of the future in $\mathcal{FL}_\gamma$.

$$SendRef_{\delta\to\gamma}(f^{\alpha\to\beta}, loc) \triangleq \; Reg_\delta(f^{\alpha\to\beta}, \gamma, FR_\delta); \;\; Reg_\gamma(f^{\alpha\to\beta}, loc, \mathcal{FL}_\gamma)$$

3.2.2 Future Computed

Once the value of a future $f^{\alpha\to\beta}$ has been computed at process β, it is immediately sent to all the processes that belong to $\mathcal{FR}_\beta(f^{\alpha\to\beta})$. This will trigger chains of *SendValue* operations. Once the future value have been sent, the future recipient list is no longer useful:

$$FutureComputed_\beta(f^{\alpha\to\beta}, value) \triangleq \; \forall\, \delta \in \mathcal{FR}_\beta(f^{\alpha\to\beta}),\; SendValue_{\beta\to\delta}(f^{\alpha\to\beta}, value) \\ Clear_\beta(f^{\alpha\to\beta}, \mathcal{FR}_\beta))$$

3.2.3 Send Future Value

When a future value is received, the receiver first updates all the local references, and then sends the future value to all the processes to which it had forwarded the future (the processes in its $\mathcal{FR}$ list). The operation is recursive, because the destination process of *SendValue* may also need to update further futures. This operation can potentially trigger the $SendRef$ operation in case of nested futures. The future locations and future recipient lists for this future are not anymore needed after those step:

$$SendValue_{\delta\to\epsilon}(f^{\alpha\to\beta}, value) \triangleq \; \forall\, loc \in \mathcal{FL}_\epsilon(f^{\alpha\to\beta}),\; Update_\epsilon(loc, value), \\ Clear_\epsilon(f^{\alpha\to\beta}, \mathcal{FL}_\epsilon) \\ \forall\, \gamma \in \mathcal{FR}_\epsilon(f^{\alpha\to\beta}),\; SendValue_{\epsilon\to\gamma}(f^{\alpha\to\beta}, value), \\ Clear_\epsilon(f^{\alpha\to\beta}, \mathcal{FR}_\epsilon)$$

3.3 Eager Message-based Strategy

In eager message-based strategy, the process β, computing the future value, is responsible for updating all nodes which receive a future. Opposed to forward-based strategy where futures updates are performed in a distributed manner, here all updates are performed by same process β (home) in a centralized manner. Whenever, a process δ forwards a future to another process γ, it sends a message *SendRegReq* to the home process β, and updates the list of future recipients $\mathcal{FR}_\beta$. $\mathcal{FR}_\beta(f^{\alpha\to\beta})$ contains the set of processes to which $f^{\alpha\to\beta}$ has been forwarded.

Figure 2 shows an example of this strategy. When A forwards the future to process B a registration message *SendRegReq* is sent from A to H, registering B in $\mathcal{FR}_H$. Similarly we have a registration message sent to H from B adding C, E, and D to $\mathcal{FR}_H$; finally we have $\mathcal{FR}_H(f^{A\to H}) = \{A, B, C, D, E, F\}$.

Once the future result is available, H uses the *SendValue* message to communicate the value to all processes in $\mathcal{FR}_H(f^{A\to H})$.

3.3.1 Send Future Reference

In the message-based strategy when a future $f^{\alpha\to\beta}$ is forwarded by a process δ to a process γ, a registration message is sent to the process that will compute the future, β.

$$SendRef_{\delta\to\gamma}(f^{\alpha\to\beta}, \gamma, loc) \triangleq \; Reg_\beta(f^{\alpha\to\beta}, \gamma, \mathcal{FR}_\beta); \;\; Reg_\gamma(f^{\alpha\to\beta}, loc, \mathcal{FL}_\gamma)$$

The registration $Reg_\beta(f^{\alpha\to\beta}, \gamma, \mathcal{FR}_\beta)$ is performed using a communication addressed to the home process β, and is called *SendRegReq* in Figure 2.

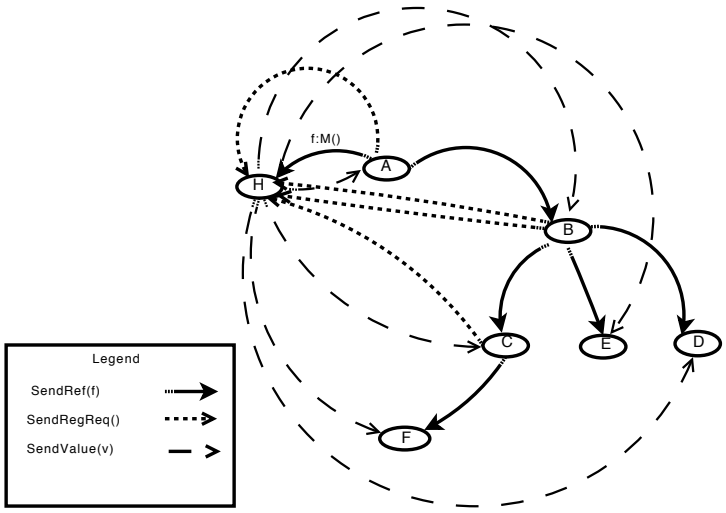

Figure 2: Future-update in eager message-based strategy

3.3.2 Future Computed

Once the execution is completed and the value is available in β, the process β sends the value to all the processes in $\mathcal{FR}_\beta(f^{\alpha\to\beta})$.

$$\begin{aligned} \mathit{FutureComputed}_\beta(f^{\alpha\to\beta}, \mathit{val}) \triangleq\ & \forall\, \delta \in \mathcal{FR}_\beta(f^{\alpha\to\beta})\ \mathit{SendValue}_{\beta\to\delta}(f^{\alpha\to\beta}, \mathit{val}); \\ & \mathit{Clear}_\beta(f^{\alpha\to\beta}, \mathcal{FR}_\beta) \end{aligned}$$

3.3.3 Send Future Value

Contrarily to forward-based strategy, there is no need to forward the future value when received, only local references are updated, and then the $\mathcal{FL}$ list can be cleared.

$$\begin{aligned} \mathit{SendValue}_{\beta\to\gamma}(f^{\alpha\to\beta}, \mathit{val}) \triangleq\ & \forall\, \mathit{loc} \in \mathcal{FL}_\gamma(f^{\alpha\to\beta})\ \mathit{Update}_\gamma(\mathit{loc}, \mathit{val}); \\ & \mathit{Clear}_\gamma(f^{\alpha\to\beta}, \mathcal{FL}_\gamma) \end{aligned}$$

The received future value may contain other futures as well. In this case, it can potentially trigger the send future reference operation.

3.4 Lazy Message-based Strategy

The lazy strategy differs from the eager strategies in the sense that future values are only transmitted when absolutely required. When a process accesses a unresolved future, the access triggers the update. This strategy is somewhat similar to message-based strategy except the futures are updated only when and if necessary. In addition, each process now needs to store all the future values that it has computed. For this, we introduce another list, $\mathcal{FV}$ that stores these values: $\mathcal{FV}\colon \mathcal{F} \mapsto \mathcal{P}(\mathit{Value})$. $\mathcal{FV}_\beta(f^{\alpha\to\beta})$, if defined, contains a singleton, which is the future value of $f^{\alpha\to\beta}$.

Compared to Figure 2, in the lazy strategy only the processes that require the future value register in $\mathcal{FR}_H$, $\mathcal{FR}_H(f^{A\to H}) = \{C, D\}$ if only C and D access the future. When the result is available, H communicates it to processes in $\mathcal{FR}_H(f^{A\to H})$. In addition, the value is stored in $\mathcal{FV}_H(f^{A\to H})$. If the future value is required later, it will be retrieved from $\mathcal{FV}_H(f^{A\to H})$.

3.4.1 Send future reference

This strategy does not require registration with home process when forwarding a future. Incoming futures are registered in FL_γ on the receiver. Once the value is received, all local references can be updated.

$$SendRef_{\delta\to\gamma}(f^{\alpha\to\beta},\gamma,loc) \triangleq Reg_\gamma(f^{\alpha\to\beta},loc,\mathcal{FL}_\gamma)$$

3.4.2 Wait-by necessity

Wait-by-necessity is triggered when the process tries to access the value of the future. We register the waiting process at β:

If the future has already been computed by β, the value is transmitted immediately. Otherwise, the request is added to the Future receivers list of β.

$$SendRegReq_{\gamma\to\beta}(f^{\alpha\to\beta},\gamma) \triangleq \begin{cases} SendValue_{\beta\to\gamma}(f^{\alpha\to\beta},val) & \text{if } \mathcal{FV}_\beta(f^{\alpha\to\beta})=\{val\} \\ Reg_\beta(f^{\alpha\to\beta},\gamma,\mathcal{FR}_\beta) & \text{if } f^{\alpha\to\beta}\notin dom(\mathcal{FV}_\beta) \end{cases}$$

3.4.3 Future Computed

When a result is computed, the value is stored in the future value list. Moreover, if there are pending requests for the value, then the value is sent to all the awaiting processes.

$$\begin{aligned} FutureComputed_\beta(f^{\alpha\to\beta},val) \triangleq\ & \forall\,\delta\in\mathcal{FR}_\beta(f^{\alpha\to\beta})SendValue_{\beta\to\delta}(f^{\alpha\to\beta},val) \\ & Clear_\beta(f^{\alpha\to\beta},\mathcal{FR}_\beta);\ \ Reg_\beta(f^{\alpha\to\beta},val,\mathcal{FV}_\beta) \end{aligned}$$

3.4.4 Send Future Value

The $SendValue$ operation is the same as for the eager message-based strategy:

$$\begin{aligned} SendValue_{\beta\to\gamma}(f^{\alpha\to\beta},val) \triangleq\ & \forall\, loc\in\mathcal{FL}_\gamma(f^{\alpha\to\beta})\ Update_\gamma(loc,val); \\ & Clear_\gamma(f^{\alpha\to\beta},\mathcal{FL}_\gamma) \end{aligned}$$

3.5 Cost Analysis of Update Strategies

In this section, we present a simple model for analyzing the cost of updating futures using different strategies. For the purpose of our analysis, we assume that futures are forwarded over a simple tree like configuration of activities. The cost analysis focuses on the time necessary for updates and does not consider the computation time which is too much application dependent. We do not measure the time necessary to require a result because this measure has a different meaning for each strategy. Thus, for message-based strategies, we suppose that, all the registration requests have been received when the result is computed. Our analysis focuses on a tree T consisting of N activities; depth of T is $D(T)$. We assume that all nodes have the same degree d. If the degree is not constant, d is the average degree of the nodes in T, and the cost is approximated. At each node, the maximum number of concurrent updates a process can perform is given by k (size of thread pool). Additionally, for lazy strategy, only l processes (nodes) out of N make use of future values. In order to update future values on various nodes in T, values must be serialized-deserialized at appropriate nodes. The time spent in serializing-deserializing for one transfer is denoted by t_s, while t_f is the time required for transferring the serialized

result. Additionally, t_r is the time for registering a process as a future recipient (for the lazy strategy). Using these notations, we aim to approximate: a) the total number of messages needed to update a future, b) The total time needed to update a given future $f^{\alpha \to \beta}$ at all N processes, c) T_w: the time for a given node/process γ to receive a result.

The following table presents a summary of the cost evaluation.

Variable	*Eager forward-based*	*Eager Message-based*	*Lazy*
Number of messages	N	$2 * N$	$2 * l$
Time to update all futures	$D(T) * (t_f + t_s) * \lceil d/k \rceil$	$\lceil N/k \rceil * t_f + t_s$	Not Applicable
Time to update a future at a given node	$t_f + t_s \leq T_w \leq D(\gamma) * (t_f + t_s) * \lceil d/k \rceil$	$t_f + t_s \leq T_w \leq \lceil N/k \rceil * t_f + t_s$	$t_s + t_f + t_r \leq T_w \leq \lceil l/k \rceil * t_f + t_s + t_r$.

In eager-forward strategy the responsibility of future update is distributed among all intermediate nodes. This can be a important consideration in environments where the bandwidth available is limited. On the other hand, this implies that future update time is dependent on the number of intermediate nodes that must be traversed. Each intermediate node requires serialization-deserialization. As can be observed from the model, the value of t_s plays a important role. In case of huge data sizes, t_f can become significantly large.

In message-based strategy the responsibility to update the future values is centralized at computing node. This can potentially over-load the process and available bandwidth. On the other hand, all updates are carried out by a single process, thus results need to be serialized only once if group communication mechanisms are employed. Also, the update time is independent from the location of the node, and all nodes receive future update in a relatively constant time. This can be an important consideration in scenarios where the t_s is relatively large and the depth of the tree is significant.

For lazy strategy, the registration requests can arrive at any time, before or after future value has been computed. The main drawback of this strategy is that, since registration is performed only when needed, the access to a future necessarily waits for one registration request plus the time to update a future. This introduces additional delay compared to the approximation above, and to the eager strategies. Lazy strategy can greatly reduce the number of messages exchanged. This can be a benefit in environments where network charge is an important consideration, or when future references spread but only a few nodes need the value. In counterpart, this strategy is costly in memory because future values must be stored indefinitely at the computing node.

4 Experimental Evaluation

We conducted an experimentation with a real system in order to validate the cost estimation proposed in the previous section. To this end, we adopted ProActive version 3.90. ProActive is based on the notion of active objects, abstracting processes with a unique thread and message queue. We used a cluster of 11 nodes equipped with Intel(R) Xeon(TM) CPUs at 2.80GHz with 1 GB RAM running Linux kernel 2.6.9. The cluster nodes are connected via

a Gigabit Ethernet link. To measure the various parameters of interest, we deployed an application featuring a tree topology where each node is an active object. For the scope of the analysis, we kept the number of nodes making strict operations constant. In addition, only the leaf nodes of the tree make use of future values. It should be noted that we have implemented the message-based approaches using java RMI, instead of a multicast api; this affects the performance of message-based and lazy strategy.

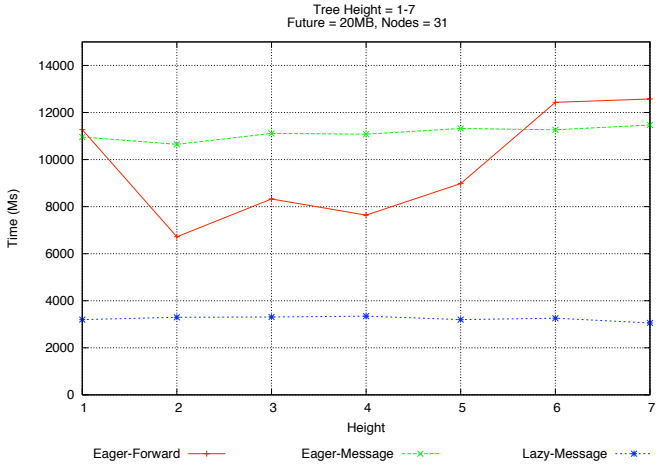

Figure 3: Comparison of strategies for a tree configuration

The graph in Figure 3 compares the time needed to update futures for the evaluated strategies. Experiments are realized over trees of varying heights. Lazy strategy takes less time to update the futures since much less updates have to be made than for the two eager strategies. As expressed in Section 3.5 the experience shows that update time required for lazy and eager message-based strategies is roughly independent of the height of the tree. Eager-forward based strategy can take advantage of concurrent updates. On the other hand, it also gets more time to reach the bottom of high trees as shown by the shape of the graph. As the height of the tree increases, overheads increases due to time spent at intermediate nodes. As a result, at height 7, the time needed for updates is higher. Note that for height 1, both eager strategies perform in a similar way because in that case both algorithms are roughly identical.

Figure 4 shows the time necessary to update a future along a simple chain of processes. Time taken by the lazy strategy is again constant and is very small because only one update is made (for the leaf node). It can be easily observed from the graph that forward-based and message-based strategies scale in a linear manner. There is no parallelization of the updates, neither for the forward-based strategy, nor for the message-based (as it is implemented in a single threaded manner). Future updates in eager forward-based strategy go through a number of intermediate steps before arriving at the last node in the chain. This introduces and additional delay for forward-based strategy. In message-based strategies, all updates are performed by same node in single step. Thus the update time is relatively constant.

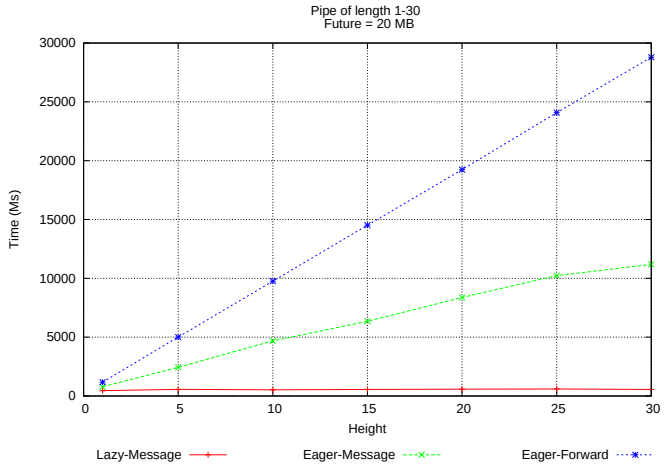

Figure 4: Comparison of strategies for a pipe configuration

5 Conclusion

This report presented a study of the three main strategies for updating first class futures. We build upon the work presented in [3, 11] to model and evaluate each of the protocols, along with experimental results. Our main contributions are:

Semi-formal event-like notation. We present and use a general (language independent) notation for modeling future update strategies. Consequently, other frameworks involving first class futures can directly benefit from our work.

Cost analysis of the strategies. For better understanding of the strategies and the relative costs (in terms of number of messages and time) involved, we present a simplified cost analysis of the protocols. This helps in understanding which strategy is more suitable for a given application.

Experimental results. We implemented the different strategies in the ProActive middleware and experimentally verified the results of our analysis.

We hope this article will help answering to the non-trivial question: "Which is the best future update strategy"? There is no single *best* strategy, rather the strategy should be adopted based on the application requirements, to summarize:

- *Eager forward-based strategy* is more suitable for scenarios where the number of intermediate nodes is relatively small and the future value is not too big. Also, the distributed nature of future updates results in less overloading at any specific node.

- *Eager message-based strategy* is more adapted for process chains since it ensures that all updates are made in relatively constant time. Due to its centralized nature, it may require more bandwidth and resources at the process that computes the future.

- *Lazy strategy* is better suited for cases where the number of processes that require future value is significantly less than total number of processes. Considerable savings in network load can be achieved but this has to be balanced against the additional delay inherent in the design of lazy approach. Also, all computed results have to be stored which requires more memory resources.

References

[1] D. Caromel, C. Delbé, A. di Costanzo, and M. Leyton. ProActive: an integrated platform for programming and running applications on grids and P2P systems. *Computational Methods in Science and Technology*, 12(1):69–77, 2006.

[2] D. Caromel, L. Henrio, and B. Serpette. Asynchronous and deterministic objects. In *Proceedings of the 31st ACM SIGPLAN-SIGACT symposium on Principles of programming languages*, pages 123–134. ACM Press, 2004.

[3] Denis Caromel and Ludovic Henrio. *A Theory of Distributed Object*. Springer-Verlag, 2005.

[4] Frank S. de Boer, Dave Clarke, and Einar Broch Johnsen. A complete guide to the future. In *ESOP*, pages 316–330, 2007.

[5] Jessie Dedecker, Tom Van Cutsem, Stijn Mostinckx, Theo D'Hondt, and Wolfgang De Meuter. Ambient-oriented programming in ambienttalk. In Dave Thomas, editor, *ECOOP*, volume 4067 of *Lecture Notes in Computer Science*, pages 230–254. Springer, 2006.

[6] Cormac Flanagan and Matthias Felleisen. The semantics of future and an application. *Journal of Functional Programming*, 9(1):1–31, 1999.

[7] Robert H. Halstead, Jr. Multilisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems (TOPLAS)*, 7(4):501–538, 1985.

[8] Einar Broch Johnsen and Olaf Owe. An asynchronous communication model for distributed concurrent objects. In *SEFM '04: Proceedings of the Software Engineering and Formal Methods*, pages 188–197, Washington, DC, USA, 2004. IEEE Computer Society.

[9] Einar Broch Johnsen, Olaf Owe, and Ingrid Chieh Yu. Creol: a type-safe object-oriented model for distributed concurrent systems. *Theor. Comput. Sci.*, 365(1):23–66, 2006.

[10] Joachim Niehren, Jan Schwinghammer, and Gert Smolka. A concurrent lambda calculus with futures. *Theoretical Computer Science*, 364(3):338–356, November 2006.

[11] Nadia Ranaldo and Eugenio Zimeo. Analysis of different future objects update strategies in proactive. In *IPDPS 2007: Parallel and Distributed Processing Symposium, IEEE International*, pages 23–66, 2007.

[12] Kenjiro Taura, Satoshi Matsuoka, and Akinori Yonezawa. Abcl/f: A future-based polymorphic typed concurrent object-oriented language - its design and implementation. In *DIMACS '94*, volume 18, 1994.

[13] G. Tretola and E. Zimeo. Extending semantics of web services to support asynchronous invocation and continuation. In *Proceedings of the IEEE International Conference on Web Services (ICWS)*, pages 208–215, 2007.

[14] G. Tretola and E. Zimeo. Activity pre-scheduling for run-time optimisation of grid workflows. *Journal of Systems Architecture*, 54(9), 2008.

[15] Adam Welc, Suresh Jagannathan, and Antony Hosking. Safe futures for java. *SIGPLAN Not.*, 40(10):439–453, 2005.

[16] Akinori Yonezawa, Etsuya Shibayama, Toshihiro Takada, and Yasuaki Honda. Modelling and programming in an object-oriented concurrent language ABCL/1. In A. Yonezawa and M. Tokoro, editors, *Object-Oriented Concurrent Programming*. MIT Press, 1987.

Centre de recherche INRIA Sophia Antipolis – Méditerranée
2004, route des Lucioles - BP 93 - 06902 Sophia Antipolis Cedex (France)

Centre de recherche INRIA Bordeaux – Sud Ouest : Domaine Universitaire - 351, cours de la Libération - 33405 Talence Cedex
Centre de recherche INRIA Grenoble – Rhône-Alpes : 655, avenue de l'Europe - 38334 Montbonnot Saint-Ismier
Centre de recherche INRIA Lille – Nord Europe : Parc Scientifique de la Haute Borne - 40, avenue Halley - 59650 Villeneuve d'Ascq
Centre de recherche INRIA Nancy – Grand Est : LORIA, Technopôle de Nancy-Brabois - Campus scientifique
615, rue du Jardin Botanique - BP 101 - 54602 Villers-lès-Nancy Cedex
Centre de recherche INRIA Paris – Rocquencourt : Domaine de Voluceau - Rocquencourt - BP 105 - 78153 Le Chesnay Cedex
Centre de recherche INRIA Rennes – Bretagne Atlantique : IRISA, Campus universitaire de Beaulieu - 35042 Rennes Cedex
Centre de recherche INRIA Saclay – Île-de-France : Parc Orsay Université - ZAC des Vignes : 4, rue Jacques Monod - 91893 Orsay Cedex

Éditeur
INRIA - Domaine de Voluceau - Rocquencourt, BP 105 - 78153 Le Chesnay Cedex (France)
http://www.inria.fr
ISSN 0249-6399